

Artificial Intelligence

Lecture 1

Dr. Yosser ATASSI

Textbook: Stuart Russell & Peter Norvig "Artificial Intelligence: A Modern Approach" Prentice Hall Series in Artificial Intelligence.

Syllabus

I Artificial Intelligence

- 1 Introduction
- 2 Intelligent Agents

II Problem-solving

- 3 Solving Problems by Searching
- 4. Informed Search and Exploration
- 5. Constraint Satisfaction Problems
- 6. Adversarial Search

III Knowledge and reasoning

- 7. Logical Agents
- 8. First-Order Logic
- 9. Inference in First-Order Logic
- 10. Knowledge Representation

Introduction

- ❖ What is AI
- ❖ AI history
- ❖ Approaches to AI


What is AI?

Systems that think like humans	Systems that think rationally
Systems that act like humans	Systems that act rationally

Acting humanly: The Turing test

Turing (1950) "Computing machinery and intelligence":

- ◇ "Can machines think?" → "Can machines behave intelligently?"
- ◇ Operational test for intelligent behavior: the Imitation Game



- ◇ Predicted that by 2000, a machine might have a 30% chance of fooling a lay person for 5 minutes
- ◇ Anticipated all major arguments against AI in following 50 years
- ◇ Suggested major components of AI: knowledge, reasoning, language understanding, learning

Problem: Turing test is not **reproducible**, **constructive**, or amenable to **mathematical analysis**

Thinking humanly

- **Cognitive science**: the brain as an information processing machine
 - Requires scientific theories of how the brain works
- How to understand cognition as a computational process?
 - **Introspection**: try to think about how we think
 - **Predict** and **test** behavior of human subjects
 - **Image** the brain, **examine** neurological data
- The latter two methodologies are the domains of cognitive science and cognitive neuroscience

Thinking rationally: Laws of Thought

Normative (or prescriptive) rather than descriptive

Aristotle: what are correct arguments/thought processes?

Several Greek schools developed various forms of logic:

notation and **rules of derivation** for thoughts;
may or may not have proceeded to the idea of mechanization

Direct line through mathematics and philosophy to modern AI

Problems:

- 1) Not all intelligent behavior is mediated by logical deliberation
- 2) What is the purpose of thinking? What thoughts **should** I have out of all the thoughts (logical or otherwise) that I **could** have?

Acting rationally

Rational behavior: doing the right thing

The right thing: that which is expected to maximize goal achievement, given the available information

Doesn't necessarily involve thinking—e.g., blinking reflex—but thinking should be in the service of rational action

Aristotle (Nicomachean Ethics):

Every art and every inquiry, and similarly every action and pursuit, is thought to aim at some good

Rational agents

An **agent** is an entity that perceives and acts

This course is about designing **rational agents**

Abstractly, an agent is a function from percept histories to actions:

$$f : \mathcal{P}^* \rightarrow \mathcal{A}$$

For any given class of environments and tasks, we seek the agent (or class of agents) with the best performance

Caveat: **computational limitations make perfect rationality unachievable**

→ design best **program** for given machine resources

AI prehistory

Philosophy	logic, methods of reasoning mind as physical system foundations of learning, language, rationality
Mathematics	formal representation and proof algorithms, computation, (un)decidability, (in)tractability probability
Psychology	adaptation phenomena of perception and motor control experimental techniques (psychophysics, etc.)
Economics	formal theory of rational decisions
Linguistics	knowledge representation grammar
Neuroscience	plastic physical substrate for mental activity
Control theory	homeostatic systems, stability simple optimal agent designs

Potted history of AI

- 1943 McCulloch & Pitts: Boolean circuit model of brain
- 1950 Turing's "Computing Machinery and Intelligence"
- 1952–69 Look, Ma, no hands!
- 1950s Early AI programs, including Samuel's checkers program, Newell & Simon's Logic Theorist, Gelernter's Geometry Engine
- 1956 Dartmouth meeting: "Artificial Intelligence" adopted
- 1965 Robinson's complete algorithm for logical reasoning
- 1966–74 AI discovers computational complexity
Neural network research almost disappears
- 1969–79 Early development of knowledge-based systems
- 1980–88 Expert systems industry booms
- 1988–93 Expert systems industry busts: "AI Winter"
- 1985–95 Neural networks return to popularity
- 1988– Resurgence of probability; general increase in technical depth
"Nouvelle AI": ALife, GAs, soft computing
- 1995– Agents, agents, everywhere . . .
- 2003– Human-level AI back on the agenda

Approaches to AI

- ❖ Learning
- ❖ Rule-Based Systems
- ❖ Search
- ❖ Planning
- ❖ Recommendation Systems
- ❖ Robotics
- ❖ Agents

Learning


- ❖ Explanation

- ❖ Discovery
- ❖ Data Mining

- ❖ No Explanation



- ❖ Neural Nets
- ❖ Case Based Reasoning

Learning: Explanation



Learning: No Explanation

- Neural nets



Rule-Based Systems

- Logic Languages
 - Prolog, Lisp
- Knowledge bases
- Inference engines

Rule-Based Languages: Prolog

Father(ahmad, ali).

Father(ali, saeed).

Father(ali, samar).

Father(ali, aya).

Male(ahmad).

Male(ali), Male (saeed).

Female(samar).

Female(aya).

Son(X,Y) $\leftarrow$ Father(Y,X), Male(X).


Daughter(X,Y) $\leftarrow$ Father(Y,X), Female(X).

Son(saeed, ali)?

Search


- “All AI is search”
 - Game theory
 - Problem spaces
- Every problem is a “virtual” tree of all possible (successful or unsuccessful) solutions.
- The trick is to find an efficient search strategy.

Search: Game Theory



$$9! + 1 = 362,880$$

Natural Language Recognition



Artificial Intelligence Systems

